

Lyapunov Exponents for Diversity in Differentiable Games

Jonathan Lorraine
University of Toronto

Paul Vicol
University of Toronto

Jack Parker-Holder
University of Oxford

Tal Kachman
Radboud University

Luke Metz
Google Research, Brain Team

Jakob Foerster
University of Oxford

ABSTRACT

Ridge Rider (RR) is an algorithm for finding diverse solutions to optimization problems by following eigenvectors of the Hessian (“ridges”). RR is designed for conservative gradient systems (i.e., settings involving a single loss function), where it branches at saddles – easy-to-find bifurcation points. We generalize this idea to non-conservative, multi-agent gradient systems by proposing a method – denoted Generalized Ridge Rider (GRR) – for finding arbitrary bifurcation points. We give theoretical motivation for our method by leveraging machinery from the field of dynamical systems. We construct novel toy problems where we can visualize new phenomena while giving insight into high-dimensional problems of interest. Finally, we empirically evaluate our method by finding diverse solutions in the iterated prisoners’ dilemma and relevant machine learning problems including generative adversarial networks.

KEYWORDS

Multi-agent Learning; Machine Learning; Deep Learning

ACM Reference Format: Jonathan Lorraine, Paul Vicol, Jack Parker-Holder, Tal Kachman, Luke Metz, Jakob Foerster. 2022. Lyapunov Exponents for Diversity in Differentiable Games. In *Proc. of the 21st Int’l Conf. on Autonomous Agents and Multiagent Systems (AAMAS ’22), Online, May 9–13, 2022, IFAAMAS, 11 pgs.*

1 INTRODUCTION

In machine learning it is often useful to select particular solutions with desirable properties that an arbitrary (global or local) minimum might not have. For example, finding solutions in image classification using shapes which generalize more effectively than textures. Important instances of this in single-objective minimization are seeking solutions that generalize to unseen data in supervised learning [1, 2], in policy optimization [3], and generative models [4]. Many real-world systems are not so simple and instead involve multiple agents each of which uses a different subset of parameters to minimize their own objective. Some examples are generative adversarial networks (GANs) [5, 6], actor-critic models [6], curriculum learning [7–10], hyperparameter optimization [11–17], adversarial examples [18, 19], learning models [20–22], domain adversarial adaptation [23], neural architecture search [24–28], multi-agent settings [29] and meta-learning [30–32]. In these settings, the aim is to find one equilibrium (of potentially many equilibria) where agents exhibit some desired behavior.

Correspondence to: Jonathan Lorraine lorraine@cs.toronto.edu *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online. © 2022 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.*

For example, in the iterated prisoners’ dilemma (Sec. 5), solutions favoring reciprocity over unconditional defection result in higher returns for all agents. In GANs, solutions often generate a subset of the modes from the target distribution [33], and in Hanabi, some solutions coordinate far better with humans [34]. Existing methods often find solutions in small subspaces – even after many random restarts, as shown in Table 1. By finding a diverse set of equilibria in these games, we may be able to (a) find solutions with a better joint outcome, (b) develop stronger generative models in adversarial learning, or (c) find solutions that coordinate better with humans.

Recently, Ridge Rider (RR) [35] proposed a general method for finding diverse solutions in *single-objective* optimization. RR is a branching tree search, which starts at a stationary point and then follows different *eigenvectors of the Hessian* (“ridges”) with negative eigenvalues, moving downhill from saddle point to saddle point. In settings where multiple agents each have their own objective (i.e., games), the relevant generalization of the Hessian – the *game Hessian* [36] in Eq. 6 – is not symmetric. Thus, in general, the game Hessian has complex eigenvalues (EVal) with associated complex eigenvectors (EVecs), making RR not directly applicable.

In this paper, we generalize RR to multi-agent settings by leveraging machinery from *dynamical systems*. We connect RR with methods for finding *bifurcation* points, i.e. points where small changes in the initial parameters lead to very different learning dynamics and optimization outcomes. We propose novel metrics, inspired by *Lyapunov exponents* [37] that measure how quickly learning trajectories separate. Our contributions include:

- Connections between finding diverse solutions and Lyapunov exponents, allowing us to leverage the broad body of work in dynamical systems.
- Proposing a method, Generalized Ridge-Rider (GRR), that scales to high-dimensional differentiable games.
- Compared to existing methods, GRR finds diverse solutions in the IPD, spanning cooperation, defection and reciprocity.
- Lastly, we present larger-scale experiments on GANs – a model class of high interest to the ML community.

2 BACKGROUND

App. Table 4 summarizes our notation. Consider the single-objective optimization problem:

$$\theta^* := \arg \min_{\theta} \mathcal{L}(\theta) \quad (1)$$

We denote the gradient of the loss at parameters θ^j by $g^j := g(\theta^j) := \nabla_{\theta} \mathcal{L}(\theta)|_{\theta^j}$. We can locally minimize the loss \mathcal{L} using gradient descent with step size α :

$$\theta^{j+1} = \theta^j - \alpha g^j \quad (2)$$

Due to the potential non-convexity of the \mathcal{L} , multiple stationary points can exist, and gradient descent will only find a particular solution based on the initialization θ^0 .

2.1 Ridge Rider

Ridge Rider (RR) [35] finds diverse solutions in single-objective minimization problems. The method first finds a saddle point, either analytically, e.g. in tabular reinforcement learning, or by minimizing the gradient norm.

Then, RR branches the optimization procedure following different directions (or “ridges”) given by the EVecs of the Hessian $\mathcal{H} = \nabla_{\theta} \mathbf{g} = \nabla_{\theta} (\nabla_{\theta} \mathcal{L})$. Full computation of the eigendecomposition of \mathcal{H} , i.e. its EVecs and EVals, is often prohibitively expensive; however, we can efficiently access a subset of the eigenspaces via Hessian-vector products $\mathcal{H}v = \nabla_{\theta} ((\nabla_{\theta} \mathcal{L})v)$ [38–41].

2.2 Optimization in Games

Instead of simply optimizing a single loss, optimization in games involves multiple agents, each with a loss function that can depend on other agents. For simplicity, we look at 2-player games with players (denoted by A and B) who want to minimize their loss – $\mathcal{L}_A(\theta_A, \theta_B)$ or $\mathcal{L}_B(\theta_A, \theta_B)$ – with their parameters – θ_A or θ_B .

$$\theta_A^* \in \arg \min_{\theta_A} \mathcal{L}_A(\theta_A, \theta_B^*), \theta_B^* \in \arg \min_{\theta_B} \mathcal{L}_B(\theta_A^*, \theta_B) \quad (3)$$

If \mathcal{L}_B and \mathcal{L}_A are differentiable in θ_B and θ_A , we say the game is differentiable. One of the simplest optimization methods is to find local solutions by simply following the players’ gradients, but this is unstable when the game Hessian has complex EVals [42]. Here, $\mathbf{g}_A^j := \mathbf{g}_A(\theta_A^j, \theta_B^j)$ is an estimator for $\nabla_{\theta_A} \mathcal{L}_A |_{\theta_A^j, \theta_B^j}$, and the simultaneous gradient update is:

$$\theta_A^{j+1} = \theta_A^j - \alpha \mathbf{g}_A^j, \quad \theta_B^{j+1} = \theta_B^j - \alpha \mathbf{g}_B^j \quad (\text{SimSGD})$$

We simplify notation with the concatenation of all players’ parameters (or joint-parameters) $\omega := [\theta_A, \theta_B] \in \mathbb{R}^d$ and the joint-gradient vector field $\hat{\mathbf{g}} : \mathbb{R}^d \rightarrow \mathbb{R}^d$, denoted at the j^{th} iteration by:

$$\hat{\mathbf{g}}^j := \hat{\mathbf{g}}(\omega^j) := [\mathbf{g}_A(\omega^j), \mathbf{g}_B(\omega^j)] = [\mathbf{g}_A^j, \mathbf{g}_B^j] \quad (4)$$

We write the next iterate in (SimSGD) with fixed-point operator F :

$$\omega^{j+1} = F_{SGD}(\omega^j) = \omega^j - \alpha \hat{\mathbf{g}}^j \quad (5)$$

The Jacobian of the fixed point operator F – denoted J – is useful for analysis, including bounding convergence rates near fixed points [43] and finding points where local changes to parameters may cause convergence to qualitatively different solutions [44]. The fixed point operator’s Jacobian crucially depends on the Jacobian of the joint-gradient $\hat{\mathbf{g}}$, which is called the *game Hessian* [45] because it generalizes the Hessian:

$$\hat{\mathcal{H}} := \nabla_{\omega} \hat{\mathbf{g}} = \begin{bmatrix} \nabla_{\theta_A}^2 \mathcal{L}_A & \nabla_{\theta_A} \nabla_{\theta_B} \mathcal{L}_A \\ \nabla_{\theta_B} \nabla_{\theta_A} \mathcal{L}_B & \nabla_{\theta_B}^2 \mathcal{L}_B \end{bmatrix} \quad (6)$$

$$J_{SGD} := \nabla_{\omega} F_{SGD}(\omega) = I - \alpha \hat{\mathcal{H}} \quad (7)$$

In Fig. 2 we show a game with a solution that we can only converge to by using an appropriate optimizer. Thus, we need to incorporate information about the optimizer when generalizing RR, which we do by working with the (largest) EVals/EVecs of J instead of (most negative) EVals/EVecs of $\hat{\mathcal{H}}$.

We can understand the difference between optimization with a single and multiple objectives as follows: In single-objective optimization following the gradient forms trajectories from a conservative vector field because $\hat{\mathcal{H}} = \mathcal{H}$ is the Hessian of the loss which is symmetric and has real EVals. However, in games with multiple objectives, $\hat{\mathcal{H}}$ can be non-symmetric and have complex

EVals, resulting in a non-conservative vector field from optimization, opening the door to many new phenomena.

3 METHODS FOR GENERALIZING RR

Here, we present two key contributions for generalizing RR to games. We first connect diversity in optimization to the general concept of bifurcations, places where a small change to the parameters causes a large change to the optimization trajectories. Second, we introduce Lyapunov exponents [37] and easy-to-optimize variants as a tool for finding these aforementioned bifurcations.

3.1 Connecting Diversity and Bifurcations

In dynamical systems, *bifurcations* are regions of the parameter space where small perturbations result in very different optimization trajectories, and in general a dynamical system can contain a variety of different *types of bifurcations*. In contrast, in conservative gradient vector fields, saddle points are the only relevant class of bifurcations. As a consequence, their EVecs play a key role in the shape of the phase portraits, which are geometric representations of the underlying dynamics. In particular, the negative EVecs are orthogonal to separatrices [46], boundaries between regions in our system with different dynamical behavior, thus providing directions to move in for finding different solutions. This perspective provides a novel view on RR. RR branches at saddle points, the only relevant class of bifurcation points in single loss optimization.

However, in the dynamical systems literature, many bifurcation types have been studied [37]. This inspires generalizing RR to non-conservative gradient fields (e.g. multi-agent settings) where a broad variety of bifurcations occur. See Fig. 2 for a *Hopf bifurcation* [44] or Fig. 7 for various others.

3.2 Lyapunov Exponents for Bifurcations

Using tools from dynamical systems research we look at *how* to find general bifurcation points. Our objectives are inspired by the Lyapunov exponent, which measures asymptotic separation rates of optimization trajectories for small perturbations. We propose a similar quantity, but for finite length trajectories. Given a k -step trajectory generated by our fixed point operator F – i.e., optimizer – with initialization ω_0 and Jacobian at iteration j of J^j , we measure the separation rate for an initial, normalized displacement \mathbf{d} with:

$$\hat{\lambda}_k(\omega_0, \mathbf{d}) = \frac{1}{k} \sum_{j=0}^{k-1} \gamma_j(\omega_0, \mathbf{d}), \quad (8)$$

$$\text{where } \gamma_j(\omega_0, \mathbf{d}) := \log(\mathbf{d}^T (J^j(\omega_0))^T J^j(\omega_0) \mathbf{d}) \quad (9)$$

We call γ_j the j^{th} *Lyapunov term*. When $k = 0$, the $\hat{\lambda}$ are called the *local Lyapunov exponents*, while as $k \rightarrow \infty$ these are called the (*global*) *Lyapunov exponents* [47]. For a variable k , we denote this as the k -*step or truncated Lyapunov exponent*. Fig. 1 visualizes the exponents’ calculation providing additional intuition. For notational simplicity, we suppress the dependency of J^j on ω_0 going forward.

Within a basin of attraction to a given fixed point the global Lyapunov exponent is constant [37]. Intuitively, this is because an arbitrarily high number of Lyapunov terms near the fixed point dominate the average defining the exponent in Eq. 8. This property prevents us from optimizing the global exponent using gradients, making it a poor objective for bifurcations. As such, our interest in the truncated exponent is motivated from multiple directions:




Figure 1: Visualization of the components to a Lyapunov exponent $\hat{\lambda}_k(\omega_0, d)$ described in Eqs. 8, 9, which measures how quickly trajectories separate starting at a point ω_0 in direction d . Here, the optimization trajectory iterates ω_j accumulate at a fixed point ω^* . We show two displacements $-d_1$ and $-d_2$ – resulting in separate “perturbed” trajectories shown in blue and green. We measure the separation rate between the true and perturbed trajectories at the j^{th} optimizer iteration with the Lyapunov term $\gamma_j(\omega_0, d)$. The exponent $\hat{\lambda}_k(\omega_0, d)$ is the average of the first k terms. See Fig. 6 for actual trajectories on a toy problem used in an exponent calculation.

- (1) Non-zero gradient signals for finding bifurcations
- (2) Computationally tractability
- (3) A better separation rate description for the finite trajectories used in practice

However, unlike the global exponent, the truncated version lacks theoretical results.

In more than one dimension, the k -step Lyapunov exponent is a function of the specific direction of the perturbation d [46]. We look at using the direction for maximal separation – i.e., the *max k -step Lyapunov exponent*:

$$\hat{\lambda}_k^{\max}(\omega_0) = \max_{d, \|d\|=1} \hat{\lambda}_k(\omega_0, d) \quad (10)$$

For dynamical systems with basins of attraction, common in optimization, the max exponent is largest on the boundary between basins, which motivates maximizing the max exponent to find bifurcations. The max exponent can be evaluated by finding the largest EVal of an average of the Jacobians over the optimization steps [37]:

$$J^\dagger := \frac{1}{k} \sum_{j=0}^k (J^j)^\top (J^j), \quad (11)$$

$$\hat{\lambda}_k^{\max}(\omega_0) = \max_{\lambda \in \text{Sp}(J^\dagger)} |\lambda| \quad (12)$$

Importantly, in higher dimensions one eigenvalue dominates the spectrum of J after a large number of steps [48, 49] and is thus a point of maximal exploration in our solution space.

We note some practical points for computing these exponents: when $k = 0$ the max exponent is the max eigenvalue of J^0 . As $k \rightarrow \infty$ and our fixed point operator converges to a fixed point ω^* , the max exponent is the max EVal of J at ω^* . Calculating $\hat{\lambda}_k^{\max}$ is easiest when $k = 0$ or $k \rightarrow \infty$, e.g., by power iteration on the relevant J . For intermediary k , directly using leading EVals of J^\dagger involves re-evaluating the entire optimization trajectory many times. Instead, it is often easier to work with bounds. A simple lower bound is formed by using the leading EVec at any single step, or an upper bound by using the leading EVec at each step, which are tight as $k \rightarrow \infty$ [48]. We investigate these strategies in App. Fig. 9.

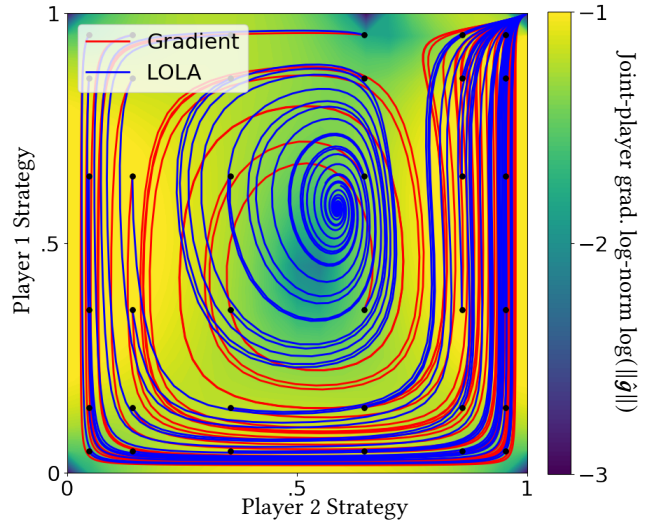


Figure 2: The phase portrait for two standard optimization algorithms on the mixed small IPD and Matching pennies problem. We show trajectories following the gradient with simSGD in red and LOLA [29] – a method for learning in games – in blue. All initializations of SimSGD only find the solution in the top right, because the center solution has imaginary EVals, while LOLA finds all solutions. For comparisons over more test problems see App. Fig. 8.

Commonly, our goal is to obtain many qualitatively different solutions from a single starting point, which motivates simultaneously optimizing the exponents corresponding to multiple different directions. Relatedly, the sum of positive global Lyapunov exponents gives an estimate of the Kolmogorov–Sinai or metric entropy by Pesin’s theorem [50]. We use this to motivate different performance metrics in Section 4.2.

4 PROPOSED ALGORITHMS

Having given an overview of the key mathematical concepts, we now present our overall algorithm. First, we introduce a general branching-tree search framework for finding diverse solutions in differentiable games. Next, we present our method – Generalized Ridge Rider (GRR) – which implements this framework using truncated Lyapunov exponents (Eq. 10) as the branching criterion. Lastly, we highlight the differences between GRR and RR.

4.1 Branching Optimization Tree Searches

Our framework is a generalized version of RR and contains the following components:

- (1) A method for finding a suitable *starting point* for our branching process - see Fig. 6.
- (2) A process for selecting *branching directions* (or perturbations) from a given branching point - see Fig. 4.
- (3) A prescription for how to *continue the optimization process* along a given branch after the initial perturbation.
- (4) A re-branching decision rule, i.e., when to go back to step (2). This was important in RR because optimizers in high-dimensional non-convex ML problems often finish at saddle points [51].
- (5) Lastly, a metric to rank the different solutions.

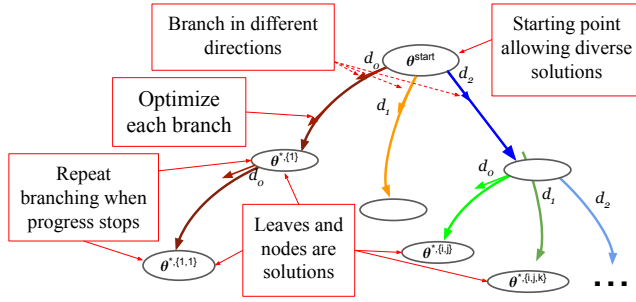


Figure 3: Visualization of branching optimization tree search. The key components are: (1) selecting the starting point, (2) creating different branches, (3) optimizing each branch, and (4) choosing when to re-branch.

We visualize this process in Fig. 3. RR is an instance of this general process, where each component is suitable for single-objective optimization. In the next section, we present another instance of this method, designed for optimization in games. We include a more detailed description of branching tree searches in App. Alg. 1, highlighting the important changes compared to RR.

4.2 Generalized Ridge Rider (GRR)

Starting point: Motivated by Section 3.2, we look at optimizing the maximal k -step Lyapunov exponent from Eq. 10 to obtain our starting point:

$$\mathcal{L}(\omega_0) = -\hat{\lambda}_k^{\max}(\omega_0) = -\max_{\mathbf{d}, \|\mathbf{d}\|=1} \hat{\lambda}_k(\omega_0, \mathbf{d}) \quad (13)$$

However, using a single exponent only guarantees trajectory separation in a single direction. If we want to branch across multiple bifurcations in different directions, we need an objective using exponents in multiple, different directions. We look at the simple objective choice summing over exponents:

$$\mathcal{L}_n^{\text{sum}}(\omega_0) = -\max_{\mathbf{d}_1, \dots, \mathbf{d}_n} \sum_{l=1}^n \hat{\lambda}_k(\omega_0, \mathbf{d}_l), \quad (14)$$

$$\text{such that } \|\mathbf{d}_l\| = 1, \mathbf{d}_l^\top \mathbf{d}_m = 0 \text{ for all } l, m \in 1, \dots, n, l \neq m \quad (15)$$


Intuitively, the constraint guarantees we have different directions to separate in, by making them orthogonal. It is straightforward to evaluate this objective, by evaluating the top- n EVal of the matrix from Eq. 11. More generally, convex functions of the k -step exponents in different directions form reasonable objectives that are more amenable to optimization. Specifically, we also look at:

$$\mathcal{L}_n^{\text{min}}(\omega_0) = -\max_{\mathbf{d}_1, \dots, \mathbf{d}_n} \min_{l=1, \dots, n} \hat{\lambda}_k(\omega_0, \mathbf{d}_l) \quad (16)$$

$$\text{such that } \|\mathbf{d}_l\| = 1, \mathbf{d}_l^\top \mathbf{d}_m = 0 \text{ for all } l, m \in 1, \dots, n, l \neq m \quad (17)$$

Branching the parameter optimization: We must choose what direction to branch in; our procedure for evaluating Lyapunov exponent objectives creates natural candidates. Specifically, evaluating the max exponent involves finding the direction maximizing trajectory separation, which we re-use for branching. Notably, this is the most negative EVal of the Hessian if we start at a saddle point and use SGD when calculating the trajectories, generalizing the choice from RR. For each direction, we can move in both a positive and negative direction, giving two branches.

GRR – Our Method – at a Saddle Bifurcation



GRR at a Hopf Bifurcation

Figure 4: We show branching at different types of bifurcations, obtained by optimizing a Lyapunov exponent as shown in Fig. 6. In each setup, we have two EVecs and branch in opposite directions, giving four paths, displayed in different colors. Steps with the eigenvector have magenta circles marking boundaries. Top: In the small IPD, finding, then branching at a saddle – where the joint-player grad. $\log\text{-norm } \log(\|\hat{g}\|)$ is 0 – allows us to find defect-defect and tit-for-tat solutions. Bottom: In the Mixed Problem of Small IPD and Matching Pennies, branching at the Hopf bifurcation allows us to find both solutions. Here, there are no saddle points near the bifurcation, so RR’s starting point does not allow branching to find both solutions.

Also, we must choose how far to move in the directions. If we move too far, we may leap into entirely different parts of the parameter space – e.g., missing interesting regions and recovering similar solution modes. If we are not exactly at a bifurcation – only near it – then we may need to move some minimum distance to cross the separatrix and find a new solution. We look at two simple strategies to move sufficiently far. First, we try taking a single step with the normalized exponent direction scaled according to the exponent. Second, we look at taking small steps in the exponent direction until the alignment with the joint-gradient flips, which generalizes RR’s “riding a ridge” (following an EVec of the Hessian) while it is a descent direction.

Optimizing each branch: For optimization in games, the stability properties of solutions can crucially depend on optimizer choice [52]. One should choose an optimizer suited to the problem. In our experiments, we use Learning with Opponent Learning Awareness (LOLA) [29] which can converge to periodic solutions and is attracted to high-welfare solutions in the IPD. In Section 6.1.1 we contrast finding diverse solutions using LOLA with simultaneous SGD (simSGD) – a method that works well for single-objective optimization, but cannot find periodic solutions. App. A summarizes other optimizer choices.

Re-branching: In single-objective optimization in ML, our optimizer often finishes at a high-dimensional saddle [51], which makes re-branching important. Specifically, we can re-branch at the saddle in negative EVec directions to try to find critical points with less negative EVecs. In our setup, we are interested in re-branching if our optimizer finishes at a point where EVals of the Jacobian of the fixed point operator J are greater than 1. These are directions where our optimizer will continue moving the parameters. [53] observed EVals of J larger than 1 at the end of GAN training, indicating that we may want to re-branch for games in machine learning.

4.3 Comparing GRR and RR

RR is a branching optimization search specifically for single-objective optimization – which is less general than optimization in games – so it can outperform GRR. E.g., non-conservative systems have more bifurcation types than conservative ones. If we are only concerned with saddle bifurcations, we can just find a saddle stationary point by minimizing the gradient norm. We know that this (relatively) easy-to-find stationary point lies on the separatrix. However, Hopf bifurcations are not necessarily near stationary points. Thus optimizing gradient norms does not work in general, while optimizing a Lyapunov exponent does (Fig. 2).

While it might be overkill to find a separatrix with Lyapunov exponents in a single-objective setting, we take some lessons from GRR back to RR. It is useful to view RR as a method for finding bifurcations and branching across them. This motivates ways to sort between different stationary points to start at – an open problem from RR. For example, using the point with the largest Kolmogorov-Sinai entropy [50]. At stationary points, this is simply the (negative) sum of negative EVals. Another limitation of RR is effectively estimating the most negative EVals of the Hessian. It is often simpler – in computation and implementation – to estimate the leading EVals of the Jacobian of the fixed point operator J instead of the most negative eigenvalues of the Hessian \hat{H} . In Section 6.2.3 we show that our method saves Hessian-vector product evaluations when estimating EVecs in setups from RR.

5 EXPERIMENTAL SETTING

We experimentally investigate GRR on a variety of problems summarized in this section and described in detail in App. C.1. We chose these as they cover different types of dynamics and contain different kinds of bifurcations. Some are standard benchmarks, while others – i.e., Random Subspaces – are novel to this work. We also summarize our gradient computation for these problems.

Matching Pennies is a simplified 2-parameter version of rock-paper-scissors. This problem’s game Hessian has purely imaginary

EVals unlike the small IPD, but only a single solution. Thus, by itself, is a poor fit for evaluating methods for a diversity of solutions, but nevertheless a useful test when probing GRR behavior.

The **Iterated Prisoners’ Dilemma (IPD)** is the discounted, infinitely iterated Prisoner’s Dilemma [54]. Each agent’s policy conditions on the actions in the prior time step, so there are 5 parameters for each agent – the probability of cooperating initially and those given both agents’ preceding actions. There are several different relevant equilibria in the IPD, including *unconditional* defection (DD), leading to the worst-case joint outcome, and *tit-for-tat* (TT), where agents initially cooperate, then copy the opponents’ action (giving a higher reward). We turn the IPD into a differentiable game by calculating the analytical expected return as a function of the joint policy of the two agents.

The **Small IPD** is a 2-parameter simplification of IPD, which has both DD and TT Nash equilibria, allowing us to visualize optimization difficulties from the full-scale IPD. However, the game Hessian has strictly real EVals, unlike the full-scale IPD.

Mixing Small IPD and Matching Pennies interpolates between the Small IPD and Matching pennies games with an interpolation factor $\tau \in [0, 1]$. This problem has two solutions – one where both players cooperate and one where both players select actions uniformly, with a Hopf bifurcation separating these.

Generative Adversarial Networks (GANs): We use a setup from [55, 36, 56], where the task is to learn a Gaussian mixture distribution using GANs. The data is sampled from a multimodal distribution to investigate the tendency to collapse on a subset of modes during training – see App. Fig. 16 for the ground truth.

Random Subspace IPD/GAN: To see how robustly we can find bifurcations with the exponents, we construct more complicated toy problems by taking higher-dimensional problems and optimizing in a random subspace. For each player, we select a random direction to optimize in, by sampling a vector \mathbf{v} with entries from $U[0, 1]$ and normalizing it. Additionally, we select a random offset \mathbf{b} from whatever an appropriate initialization is for the higher-dimensional problem. So, the first player controls x -coordinate and optimizes the loss $\mathcal{L}_A(\mathbf{v}_A x + \mathbf{b}_A, \mathbf{v}_B y + \mathbf{b}_B)$, while the second player controls the y -coordinate and optimizes $\mathcal{L}_B(\mathbf{v}_A x + \mathbf{b}_A, \mathbf{v}_B y + \mathbf{b}_B)$.

Single-objective problems: We apply our method to find bifurcations on single-objective optimization problems. There are various relevant problems in machine learning, but we focus on comparisons with RRs EVec estimation in MNIST classification.

Optimizing the starting point objective: To optimize these objectives, we use automatic differentiation libraries (like Jax [41] or PyTorch [40]) to compute gradients through methods that calculate our Lyapunov exponent-inspired objectives. The scalability of this approach depends on the implementation of our exponent calculation, which can depend on estimating the top eigenvalues of the positive semi-definite (PSD) symmetric matrix in Eq. 11. In simple settings we can differentiate through the full spectrum calculation via `jax.linalg.numpy.eigh`; we investigate this on the toy experiment in Sec. 6.1.2 and the IPD in Sec. 6.2.1. However, in ML, the matrix from Eq. 11 is typically too large for the full spectrum. Directly estimating the top EVals with an iterative method allows us to (automatically) differentiate through them. We differentiate through `jax.numpy.linalg.eigh` in Fig.6 and investigate using power iteration with Hessian-vector products in App. Fig. 9.




Figure 5: We display gradient descent optimization on the 1-step max Lyapunov exponent objective (Eq. 13) on the IPD. *Take-away:* We effectively reduce our loss and correspondingly raise the max Eval of J . *Left:* We display our loss – i.e., the negative Lyapunov exponent objective – as optimization progresses. *Middle:* We visualize the spectrum of the Jacobian of our fixed point operator in log-polar coordinates as optimization progresses. The spectrum is shown with a scatter-plot in blue, with a progressively larger alpha at each iteration. The final spectrum is shown in red. A vertical red line is shown where the EVAL norm equals 1, signifying the cutoff between (locally) convergent and divergent eigenspaces. We effectively maximize the norm of the largest EVAL. *Right:* We display the spectrum of the game Hessian. A horizontal red line is shown where the real part of the EVAL transitions from negative to positive, signifying the cutoff between (locally) convergent and divergent eigenspaces under gradient flow. Log-polar coordinates are required to see structure in the spectrum.

6 EXPERIMENTAL RESULTS

First, in Sec. 6.1 we use the diagnostic problems to demonstrate and ablate the key parts of our algorithms – i.e. optimizer choice (Sec. 6.1.1), starting point selection (Sec. 6.1.2), and branching (Sec. 6.1.3). Next, in Sec. 6.2 we scale GRR to *large-scale problem settings* by (a) demonstrating that we improve RR’s EVec estimation for neural network classifiers in Sec. 6.2.3, and (b) calculating Lyapunov exponents for GANs in Sec. 6.2.4.

6.1 Diagnostic Experiments

6.1.1 Optimizer Choice. Here, we give a system with complex EVecs showing (a) the importance of selecting a convergent optimizer in GRR, and (b) an example task where RR cannot be applied. Fig. 2 shows the phase portrait for baseline methods on our Mixed Problem. LOLA (and other game optimizers) can find both solutions, while naïvely following the gradient always finds a single solution.

6.1.2 Starting Point Selection. Fig. 6 shows the effect of optimizing the starting point for the max 10-step Lyapunov exponent on the Mixed Problem. We find gradient-based optimization can find bifurcations. Next, Fig. 9 contrasts different direction choices for the exponent calculation. We find that re-estimating the top EVecs at each iteration performs best, though the simple methods also work. App. Fig. 10 shows the max k -step exponent for multiple numbers of steps k , showing that a moderate number of steps – e.g., 10 – allows us to find bifurcations. App. Fig. 15 shows different Lyapunov exponent objectives, trying to guarantee trajectory divergence in multiple directions. We can find bifurcations while guaranteeing trajectory separation in every direction.

Impact of inner optimizer choices on bifurcation structure: App. Fig. 11 contrasts the exponents for LOLA and simSGD, showing that we find optimizer-dependent bifurcations. App. Fig. 12 investigates the impact of optimization-algorithm parameter choices on bifurcation structure. This shows that if the step size is too large, the optimizer does not converge, resulting in bifurcations between complicated limit cycle trajectories [57], and making GRR difficult to apply.

Starting points on single-objective problems: App. Fig. 13 investigates our algorithm in single-objective problems, showing that our method finds bifurcations in the same setup as RR. App. Fig. 14 shows our method on the logistic map, giving intuition for our method on a canonical example for bifurcations.

6.1.3 Branching at Bifurcations. In Fig. 4 we demonstrate branching at bifurcations to find multiple solutions to toy problems. This shows the branching process, and an explicit example where RR’s starting point does not work, but GRR’s does. The small IPD has a saddle bifurcation, while the Mixed Problem has a Hopf bifurcation.

6.1.4 A Range of Complicated Toy Problems. In Fig. 7 we look at calculating Lyapunov exponents on toy problems with more complicated bifurcation structures. We create a variety of more complex toy problems by taking the high-dimensional IPD and GAN problems and selecting a random subspace to optimize in. We can effectively highlight bifurcations in this setup.

6.2 Scaling the Results

6.2.1 Optimizing Lyapunov Exponents on IPD. Here, we investigate our ability to use gradient-based optimizers on Lyapunov exponents in the IPD. Fig. 5 shows the feasibility of using gradient descent to tune the 1-step max Lyapunov exponent. App. Fig. 18 optimizes an objective using multiple exponents, showing that we effectively optimize multiple exponents, which gives trajectory separation in multiple directions. App. Fig. 19 compares objectives using multiple exponents, showing that using the minimum of the top n exponents gives trajectory separation in all n directions, unlike the naïve choice of optimizing their sum. The sum of exponents finds solutions separating extremely fast in the top directions, while (slowly) converging in the bottom directions. In contrast, the min of the exponents does not allow convergence in the bottom directions. App. Fig. 20 compares optimizing the k -step max Lyapunov exponent for variable k , showing that we effectively minimize multi-step exponents in higher-dimensional problems if required.




Figure 6: Calculation and optimization of a max 10-step Lyapunov exponent from Eq. 10 on the mixed small IPD and Matching Pennies problem. Gradient-based optimization on this objective effectively finds the bifurcation. *Top:* We show a heatmap of the exponent, and visualize the calculation of each exponent in two regions. This involves simulating 10-step trajectories shown in red starting at the black points, then finding a direction that maximizes trajectory separation. We use this exponent to find bifurcations – in this case between the solution in the top right and the center. *Bottom:* We show optimization trajectories for gradient ascent on the exponent for a grid of initializations. For a variety of starting points, the optimization procedure finds large exponent locations (final iterate shown with red circles).

6.2.2 GRR Applied to the IPD. Here, we use our method on the IPD, where existing methods have difficulty finding diverse solutions. There are two solution modes: ones where both agents end up defecting and cooperating respectively. Table 1 compares our method to baselines of following gradients and LOLA, each run with random initializations. Our method finds both solutions modes, unlike existing approaches. We found that it was sufficient to use the max Lyapunov exponent as our objective, which only guarantees separation in 1 direction. Similarly, we found that it was sufficient to use a 1-step or local Lyapunov exponent objective, though we may require more steps to find bifurcations in other problems.




Figure 7: We show the Lyapunov exponent heatmap (as in Fig. 6) on more complicated toy problems to see how robustly we can find different bifurcations. The exponent peaks near where trajectories (shown in red) separate, showing that we find various bifurcations. See App. Fig. 17 for other sampled subspaces. Sec. 5 describes how we construct these examples by taking higher-dimensional problems and optimizing them in a random subspace. *Top:* An IPD subspace with multiple Hopf bifurcations. *Bottom:* A GAN subspace with various bifurcations.

6.2.3 Improving RR’s EVec Estimation. We investigate efficiently finding the most negative EVecs in RR by estimating the largest EVecs of the Jacobian of our fixed point operator. We measure efficiency by comparing the number of Hessian-vector product (HVP) evaluations because HVP evaluations dominate the cost of EVec estimation here. Table 2 shows how many HVP evaluations we require to reach different MNIST classifier accuracies by following EVecs. Our method can more efficiently use HVP evaluations than the RR method because we do not need to repeatedly re-estimate the most negative EVal.

We stress that this problem is not designed to train a single, strong classifier; it is easy to simply train our network by following the gradient to 100% train accuracy. This problem was selected from RR’s experiments because it requires us to accurately and efficiently estimate negative EVecs many times. A downstream use of this is training an ensemble of classifiers for generalization.

Search Strategy	Player 1 Loss \mathcal{L} [min, max]	Player 1 Strategy Distribution, [min, max]				
		$p(C_0)$	$p(C CC)$	$p(C CD)$	$p(C DC)$	$p(C DD)$
GRR: tune max Lyap + top EVec branch + simSGD	[1.000, 2.000]	[.003, .999]	[.032, .999]	[.004, .884]	[.001, .912]	[.000, .013]
GRR: tune max Lyap + top EVec branch + LOLA	[1.000, 2.000]	[.002, .999]	[.063, .993]	[.001, .910]	[.000, .922]	[.005, .103]
20 Random init + SimSGD	[1.997, 1.998]	[.043, .194]	[.142, .480]	[.041, .143]	[.055, .134]	[.001, .001]
20 Random init + LOLA	[1.000, 1.396]	[.000, 1.00]	[.093, 1.00]	[.000, .966]	[.057, 1.00]	[.000, .947]
1 Random init + top EVec branch	[2.000, 2.000]	[.001, .003]	[.027, .030]	[.003, .007]	[.008, .009]	[.000, .000]

Table 1: We show strategies for finding diverse solutions in the iterated prisoner’s dilemma (IPD). *Takeaway:* Our method finds solutions at both loss modes, while existing approaches of using random initializations, then following the gradient or using LOLA do not find diverse solutions. The IPD has two solution modes – i.e., solutions where both agents end up defecting with a loss of 2 and where both agents end up cooperating with a loss of 1 (like tit-for-tat). We assess which modes were found by showing (P)layer 1’s strategy, which is the chance of (C)ooperating given both players’ last action – ex., $p(C|DC)$ is the chance if previously P1 defected and P2 cooperated. We compare GRR flavors with just following gradients via SimSGD and LOLA [29] from random (init)ializations. We compare with 20 random inits because GRR follows at most 20 branches, and because we have 10 EVecs in either direction (+/-). GRR only branches in directions where EVals of the Jacobian of the fixed point operator are greater than 1 (i.e., trajectories locally diverge) as visualized in Fig. 5 (middle). We look at the impact of starting at an approximate bifurcation in GRR, by branching on the EVecs at a random init. If the max Lyapunov exponent is not tuned, then each branch finds the same solution.

# HVP Evaluations	MNIST Accuracy	
	Our method	Method from RR
10 000	19% (+8%)	11%
100 000	89% (+6%)	83%
1 000 000	93% (+2%)	91%

Table 2: We show how many HVP evaluations we require to reach different MNIST classifier accuracies by following EVecs, repeating the exp. in RR’s Fig. 4. This experiment is not designed to train a single strong classifier, but to test our ability to efficiently follow negative EVecs – see Sec. 6.2.3.

6.2.4 Calculating Lyapunov Exponents for GANs. Here, we investigate scaling our exponent calculations to machine learning models where the (game) Hessian is so large we cannot materialize it and can only use Hessian-vector products. Specifically, we use the GAN described in Section 5. We look at calculating our exponent for various hyperparameters and random re-starts. We evaluate the quality of using our exponent to find diverse solutions, by calculating the log-probability of samples from an ensemble of GANs from the top 5 optimization branches. Table 3 shows the mean and standard deviation (over 10 random restarts) of the max 10-step Lyapunov exponent and the resulting ensemble’s log-probability. Each GAN was trained for 10 000 updates, so evaluating each ensemble cost approximately 50 000 evaluations of both players’ gradients. In contrast, each exponent cost less than 1000 evaluations of both gradients to compute.

This shows that we effectively scale our exponent calculation to larger models of interest from machine learning, and find that a large (mean) exponent aligns with regions where we can branch to train the strongest ensemble of GANs.

Init scale, step size	Max Lyap Coeff	Ensemble log-prob
0.001, 1.0	0.952 ± 0.834	$-16\,342 \pm 817$
0.1, 1.0	6.485 ± 1.155	$-13\,691 \pm 1317$
10.0, 1.0	0.053 ± 0.128	$-46\,659 \pm 26\,793$
0.001, 0.1	0.849 ± 0.765	$-12\,321 \pm 126$
0.1, 0.1	6.571 ± 0.953	$-10\,846 \pm 256$
10.0, 0.1	-0.012 ± 0.014	$-23\,459 \pm 12\,693$

Table 3: We display the mean and standard deviation (over 10 random restarts) of the max 10-step Lyapunov exponent and the log-probability of an ensemble of 5 GANs obtained by branching in the top 5 directions at the initialization. We show that the better performing ensembles also have higher Lyapunov coefficients as well as demonstrating that our exponent calculation is scalable to larger problems. The best GANs log-prob. from the best ensemble was $-12\,861 \pm 356$, which is worse than the ensemble’s performance of $-10\,846 \pm 256$. This indicates that each GAN may be learning a different part of the data distribution (samples in App. Fig. 16).

7 CONCLUSION

In this paper we introduced Generalized Ridge Rider, an extension of the Ridge Rider algorithm to settings with multiple losses. We showed that, in these settings, a broader class of bifurcation points needs to be considered, and that GRR indeed discovers them in a variety of problems. Experimentally, we isolate each component of GRR demonstrating their effectiveness, and show that – in contrast to baseline methods – GRR obtains a diversity of qualitatively different solutions in multi-agent settings such as the iterated prisoner’s dilemma. We also provide empirical justification for our method by using tools from the dynamical systems literature, allowing us to find arbitrary bifurcations. This hints at a multitude of approaches and tools from dynamical systems, that can be used for understanding game dynamics and learning diversity.

REFERENCES

- [1] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A Wichmann, and Wieland Brendel. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. In *ICLR*, 2018.
- [2] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673, 2020.
- [3] Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. Robots that can adapt like animals. *Nature*, 521(7553):503–507, 2015.
- [4] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *ICLR*, 2020.
- [5] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2672–2680, 2014.
- [6] David Pfau and Oriol Vinyals. Connecting generative adversarial networks and actor-critic methods. *arXiv:1610.01945*, 2016.
- [7] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *International Conference on Machine Learning (ICML)*, pages 41–48, 2009.
- [8] Bowen Baker, Ingmar Kanitscheider, Todor Markov, Yi Wu, Glenn Powell, Bob McGrew, and Igor Mordatch. Emergent tool use from multi-agent autocurricula. In *ICLR*, 2019.
- [9] David Balduzzi, Marta Garnelo, Yoram Bachrach, Wojciech Czarnecki, Julien Perolat, Max Jaderberg, and Thore Graepel. Open-ended learning in symmetric zero-sum games. In *International Conference on Machine Learning (ICML)*, pages 434–443, 2019.
- [10] Sainbayar Sukhbaatar, Zeming Lin, Ilya Kostrikov, Gabriel Synnaeve, Arthur Szlam, and Rob Fergus. Intrinsic motivation and automatic curricula via asymmetric self-play. In *International Conference on Learning Representations (ICLR)*, 2018.
- [11] Justin Domke. Generic methods for optimization-based modeling. In *Artificial Intelligence and Statistics*, pages 318–326, 2012.
- [12] Dougal Maclaurin, David Duvenaud, and Ryan Adams. Gradient-based hyperparameter optimization through reversible learning. In *International Conference on Machine Learning (ICML)*, pages 2113–2122, 2015.
- [13] Jonathan Lorraine and David Duvenaud. Stochastic hyperparameter optimization through hypernetworks. *arXiv preprint arXiv:1802.09419*, 2018.
- [14] Matthew MacKay, Paul Vicol, Jon Lorraine, David Duvenaud, and Roger Grosse. Self-tuning networks: Bilevel optimization of hyperparameters using structured best-response functions. In *International Conference on Learning Representations (ICLR)*, 2019.
- [15] Aniruddh Raghu, Maithra Raghu, Simon Kornblith, David Duvenaud, and Geoffrey Hinton. Teaching with commentaries. In *International Conference on Learning Representations*, 2020.
- [16] Jonathan Lorraine, Paul Vicol, and David Duvenaud. Optimizing millions of hyperparameters by implicit differentiation. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1540–1552, 2020.
- [17] Aniruddh Raghu, Jonathan Lorraine, Simon Kornblith, Matthew McDermott, and David K Duvenaud. Meta-learning to improve pre-training. *Advances in Neural Information Processing Systems*, 34, 2021.
- [18] Avishek Joey Bose, Gauthier Gidel, Hugo Berrard, Andre Cianflone, Pascal Vincent, Simon Lacoste-Julien, and William L Hamilton. Adversarial example games. *arXiv preprint arXiv:2007.00720*, 2020.
- [19] Xiaoyong Yuan, Pan He, Qile Zhu, and Xiaolin Li. Adversarial examples: Attacks and defenses for deep learning. *IEEE Transactions on Neural Networks and Learning Systems*, 30(9):2805–2824, 2019.
- [20] Aravind Rajeswaran, Igor Mordatch, and Vikash Kumar. A game theoretic framework for model based reinforcement learning. In *International Conference on Machine Learning*, 2020.
- [21] Pierre-Luc Bacon, Florian Schäfer, Clement Gehring, Animashree Anandkumar, and Emma Brunskill. A Lagrangian method for inverse problems in reinforcement learning. *lis.csail.mit.edu/pubs*, 2019.
- [22] Evgenii Nikishin, Romina Abachi, Rishabh Agarwal, and Pierre-Luc Bacon. Control-oriented model-based reinforcement learning with implicit differentiation. *arXiv preprint arXiv:2106.03273*, 2021.
- [23] David Acuna, Guojun Zhang, Marc T Law, and Sanja Fidler. f-Domain-adversarial learning: Theory and algorithms for unsupervised domain adaptation with neural networks, 2021. URL <https://openreview.net/forum?id=WqXAKcwfZtI>.
- [24] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*, 2016.
- [25] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. In *AAAI Conference on Artificial Intelligence*, volume 33, pages 4780–4789, 2019.
- [26] Hanxiao Liu, Karen Simonyan, and Yiming Yang. DARTS: Differentiable architecture search. In *ICLR*, 2019.
- [27] Will Grathwohl, Elliot Creager, Seyed Kamyar Seyed Ghasemipour, and Richard Zemel. Gradient-based optimization of neural network architecture. In *Workshop ICLR*, 2018.
- [28] George Adam and Jonathan Lorraine. Understanding neural architecture search techniques. *arXiv preprint arXiv:1904.00438*, 2019.
- [29] Jakob Foerster, Richard Y Chen, Maruan Al-Shedivat, Shimon Whiteson, Pieter Abbeel, and Igor Mordatch. Learning with opponent-learning awareness. In *International Conference on Autonomous Agents and MultiAgent Systems (AAMA)*, pages 122–130, 2018.
- [30] Mengye Ren, Eleni Triantafillou, Sachin Ravi, Jake Snell, Kevin Swersky, Joshua B Tenenbaum, Hugo Larochelle, and Richard S Zemel. Meta-learning for semi-supervised few-shot classification. In *International Conference on Learning Representations*, 2018.
- [31] Aravind Rajeswaran, Chelsea Finn, Sham M Kakade, and Sergey Levine. Meta-learning with implicit gradients. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pages 113–124, 2019.
- [32] Mengye Ren, Eleni Triantafillou, Kuan-Chieh Wang, James Lucas, Jake Snell, Xaq Pitkow, Andreas S Tolias, and Richard Zemel. Flexible few-shot learning with contextual similarity. *arXiv preprint arXiv:2012.05895*, 2020.
- [33] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning (ICML)*, pages 214–223, 2017.
- [34] Hengyuan Hu, Adam Lerer, Alex Peysakhovich, and Jakob Foerster. “Other-Play” for zero-shot coordination. In *International Conference on Machine Learning (ICML)*, pages 4399–4410, 2020.
- [35] Jack Parker-Holder, Luke Metz, Cinjon Resnick, Hengyuan Hu, Adam Lerer, Alistair Letcher, Alexander Peysakhovich, Aldo Pacchiano, and Jakob Foerster. Ridge Rider: Finding diverse solutions by following eigenvectors of the Hessian. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 753–765, 2020.
- [36] David Balduzzi, Sebastien Racaniere, James Martens, Jakob Foerster, Karl Tuyls, and Thore Graepel. The mechanics of n-player differentiable games. In *International Conference on Machine Learning (ICML)*, pages 354–363, 2018.
- [37] Anatole Katok and Boris Hasselblatt. *Introduction to the Modern Theory of Dynamical Systems*. Cambridge University Press, 1997.
- [38] Barak A Pearlmutter. Fast exact multiplication by the Hessian. *Neural Computation*, 6(1):147–160, 1994.

- [39] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>.
- [40] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. 2017.
- [41] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: Composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- [42] James P Bailey, Gauthier Gidel, and Georgios Piliouras. Finite regret and cycles with fixed step-size via alternating gradient descent-ascent. In *Conference on Learning Theory*, pages 391–407, 2020.
- [43] Dimitri Bertsekas. *Nonlinear Programming*. Athena Scientific, 2008.
- [44] Jack K Hale and Hüseyin Koçak. *Dynamics and Bifurcations*, volume 3. Springer Science & Business Media, 2012.
- [45] Alistair Letcher, David Balduzzi, Sébastien Racaniere, James Martens, Jakob Foerster, Karl Tuyls, and Thore Graepel. Differentiable game mechanics. *The Journal of Machine Learning Research*, 20, 2019.
- [46] Michael Tabor. *Chaos and Integrability in Nonlinear Dynamics: An Introduction*. Wiley-Interscience, 1989.
- [47] Rodney CL Wolff. Local Lyapunov exponents: Looking closely at chaos. *Journal of the Royal Statistical Society*, 1992.
- [48] Vittorio Loreto, Giovanni Paladin, Michele Pasquini, and Angelo Vulpiani. Characterization of chaos in random maps. *Physica A: Statistical Mechanics and its Applications*, 232(1-2):189–200, 1996.
- [49] Tal Kachman, Shmuel Fishman, and Avy Soffer. Numerical implementation of the multiscale and averaging methods for quasi periodic systems. *Computer Physics Communications*, 221:235–245, 2017.
- [50] Yakov Borisovich Pesin. Characteristic Lyapunov exponents and smooth ergodic theory. *Uspekhi Matematicheskikh Nauk*, 1977.
- [51] Zhewei Yao, Amir Gholami, Kurt Keutzer, and Michael W Mahoney. PyHessian: Neural networks through the lens of the Hessian. In *IEEE International Conference on Big Data (Big Data)*, pages 581–590, 2020.
- [52] Gauthier Gidel, Reyhane Askari Hemmat, Mohammad Pezesghi, Rémi Le Priol, Gabriel Huang, Simon Lacoste-Julien, and Ioannis Mitliagkas. Negative momentum for improved game dynamics. In *The 22nd International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1802–1811, 2019.
- [53] Jonathan Lorraine, David Acuna, Paul Vicol, and David Duvenaud. Complex momentum for learning in games. *arXiv preprint arXiv:2102.08431*, 2021.
- [54] William Poundstone. *Prisoner’s Dilemma/John Von Neumann, Game Theory and the Puzzle of the Bomb*. Anchor Press, 1993.
- [55] Luke Metz, Ben Poole, David Pfau, and Jascha Sohl-Dickstein. Unrolled generative adversarial networks. *arXiv preprint arXiv:1611.02163*, 2016.
- [56] Alistair Letcher, Jakob Foerster, David Balduzzi, Tim Rocktäschel, and Shimon Whiteson. Stable opponent shaping in differentiable games. In *International Conference on Learning Representations*, 2018.
- [57] Steven H Strogatz. *Nonlinear Dynamics and Chaos with Student Solutions Manual: With Applications to Physics, Biology, Chemistry, and Engineering*. CRC Press, 2018.
- [58] Guido Van Rossum and Fred L Drake Jr. *Python reference manual*. Centrum voor Wiskunde en Informatica Amsterdam, 1995.
- [59] Travis E Oliphant. Python for scientific computing. *Computing in Science & Engineering*, 9(3):10–20, 2007.
- [60] Travis E Oliphant. *A guide to NumPy*, volume 1. Trelgol Publishing USA, 2006.
- [61] Stefan Van Der Walt, S Chris Colbert, and Gael Varoquaux. The NumPy array: A structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2):22–30, 2011.
- [62] Charles R Harris, K Jarrod Millman, Stéfán J van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al. Array programming with NumPy. *Nature*, 585(7825):357–362, 2020.
- [63] John D Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007.
- [64] Eric Jones, Travis Oliphant, Pearu Peterson, et al. SciPy: Open source scientific tools for Python. 2001.
- [65] Lars Kai Hansen and Peter Salamon. Neural network ensembles. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1990.
- [66] Yong Liu and Xin Yao. Ensemble learning via negative correlation. *Neural Networks*, 12(10):1399 – 1404, 1999.
- [67] Samarth Sinha, Homanga Bharadhwaj, Anirudh Goyal, Hugo Larochelle, Animesh Garg, and Florian Shkurti. Diversity inducing information bottleneck in model ensembles. *AAAI*, 2021.
- [68] Andrew Slavin Ross, Weiwei Pan, Leo A. Celi, and Finale Doshi-Velez. Ensembles of locally independent prediction models. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI*, 2020.
- [69] Zeld Mariet and Suvrit Sra. Diversity networks: Neural network compression using determinantal point processes. In *International Conference on Learning Representations (ICLR)*, May 2016.
- [70] Tianyu Pang, Kun Xu, Chao Du, Ning Chen, and Jun Zhu. Improving adversarial robustness via promoting ensemble diversity. In *International Conference on Machine Learning (ICML)*, 2019.
- [71] Joel Lehman and Kenneth O. Stanley. Exploiting open-endedness to solve problems through the search for novelty. In *Proceedings of the Eleventh International Conference on Artificial Life (Alife XI)*. MIT Press, 2008.
- [72] Benjamin Eysenbach, Abhishek Gupta, Julian Ibarz, and Sergey Levine. Diversity is all you need: Learning skills without a reward function. In *International Conference on Learning Representations (ICLR)*, 2019.
- [73] Jack Parker-Holder, Aldo Pacchiano, Krzysztof Choromanski, and Stephen Roberts. Effective diversity in population-based reinforcement learning. In *Advances in Neural Information Processing Systems (NeurIPS)*. 2020.
- [74] Justin K. Pugh, Lisa B. Soros, and Kenneth O. Stanley. Quality diversity: A new frontier for evolutionary computation. *Frontiers in Robotics and AI*, 3:40, 2016.
- [75] Guodong Zhang, Xuchan Bao, Laurent Lessard, and Roger Grosse. A unified analysis of first-order methods for smooth games via integral quadratic constraints. *Journal of Machine Learning Research*, 22(103): 1–39, 2021.
- [76] Guodong Zhang, Yuanhao Wang, Laurent Lessard, and Roger Grosse. Near-optimal local convergence of alternating gradient descent-ascent for minimax optimization.
- [77] Galina M Korpelevich. The extragradient method for finding saddle points and other problems. *Matecon*, 12:747–756, 1976.
- [78] Waïss Azizian, Ioannis Mitliagkas, Simon Lacoste-Julien, and Gauthier Gidel. A tight and unified analysis of gradient-based methods for a whole spectrum of differentiable games. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.

- [79] Sasha Rakhlin and Karthik Sridharan. Optimization, learning, and games with predictable sequences. *Advances in Neural Information Processing Systems*, 26:3066–3074, 2013.
- [80] Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. Training gans with optimism. In *International Conference on Learning Representations (ICLR 2018)*, 2018.
- [81] Guodong Zhang and Yuanhao Wang. On the suboptimality of negative momentum for minimax optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 2098–2106. PMLR, 2021.
- [82] Gauthier Gidel, Hugo Berard, Gaëtan Vignoud, Pascal Vincent, and Simon Lacoste-Julien. A variational inequality perspective on generative adversarial networks. In *International Conference on Learning Representations (ICLR)*, 2018.
- [83] Lars Mescheder, Sebastian Nowozin, and Andreas Geiger. The numerics of gans. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 1823–1833, 2017.
- [84] Eric V Mazumdar, Michael I Jordan, and S Shankar Sastry. On finding local Nash equilibria (and only local Nash equilibria) in zero-sum games. *arXiv preprint arXiv:1901.00838*, 2019.
- [85] Florian Schäfer, Anima Anandkumar, and Houman Owhadi. Competitive mirror descent. *arXiv preprint arXiv:2006.10179*, 2020.
- [86] Yuanhao Wang, Guodong Zhang, and Jimmy Ba. On solving minimax optimization locally: A follow-the-ridge approach. In *International Conference on Learning Representations (ICLR)*, 2019.
- [87] Paul Vicol, Jonathan Lorraine, David Duvenaud, and Roger Grosse. Implicit regularization in overparameterized bilevel optimization. In *ICML 2021 Beyond First Order Methods Workshop*, 2021.
- [88] Marta Garnelo, Wojciech Marian Czarnecki, Siqi Liu, Dhruva Tirumala, Junhyuk Oh, Gauthier Gidel, Hado van Hasselt, and David Balduzzi. Pick your battles: Interaction graphs as population-level objectives for strategic diversity. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems (AAMA)*, pages 1501–1503, 2021.
- [89] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575, 2019.
- [90] Yaodong Yang, Ying Wen, Jun Wang, Liheng Chen, Kun Shao, David Mguni, and Weinan Zhang. Multi-agent determinantal q-learning. In *International Conference on Machine Learning (ICML)*, 2020.
- [91] Andrei Lupu, Hengyuan Hu, and Jakob Foerster. Trajectory diversity for zero-shot coordination. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, 2021.
- [92] Peter Hart, Nils Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968.
- [93] Bruce Abramson. *The Expected-Outcome Model of Two-Player Games*. Morgan Kaufmann, 2014.
- [94] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of Go with deep neural networks and tree search. *Nature*, 2016.
- [95] Thomas M Moerland, Joost Broekens, Aske Plaat, and Catholijn M Jonker. A0c: Alpha zero in continuous action space. *arXiv preprint arXiv:1805.09613*, 2018.
- [96] Beomjoon Kim, Kyungjae Lee, Sungbin Lim, Leslie Kaelbling, and Tomás Lozano-Pérez. Monte Carlo tree search in continuous spaces using Voronoi optimistic optimization with regret bounds. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020.
- [97] Weichao Mao, Kaiqing Zhang, Qiaomin Xie, and Tamer Basar. Polyhoot: Monte-carlo planning in continuous space mdps with non-asymptotic analysis. *Advances in Neural Information Processing Systems*, 33, 2020.
- [98] Steven M LaValle et al. Rapidly-exploring random trees: A new tool for path planning. *Technical Report, Iowa State University, USA*, 1998.
- [99] Steven M LaValle and James J Kuffner Jr. Randomized kinodynamic planning. *The International Journal of Robotics Research*, 2001.
- [100] Samuel Rodriguez, Xinyu Tang, Jyh-Ming Lien, and Nancy M Amato. An obstacle-based rapidly-exploring random tree. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2006.
- [101] Gulshan Singh and Kalyanmoy Deb. Comparison of multi-modal optimization algorithms based on evolutionary algorithms. In *Proceedings of the 8th Annual Conference on Genetic and Evolutionary Computation*, pages 1305–1312, 2006.
- [102] Ka-Chun Wong. Evolutionary multimodal optimization: A short survey. In *Advances in Evolutionary Algorithms Research*. 2015.
- [103] Kai Arulkumaran, Antoine Cully, and Julian Togelius. Alphastar: An evolutionary computation perspective. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, 2019.
- [104] Jonathan Lorraine, Jack Parker-Holder, Paul Vicol, Aldo Pacchiano, Luke Metz, Tal Kachman, and Jakob Foerster. Using bifurcations for diversity in differentiable games. In *ICML 2021 Beyond First Order Methods Workshop*, 2021.
- [105] E Christopher Zeeman. Population dynamics from game theory. In *Global Theory of Dynamical Systems*, pages 471–497. Springer, 1980.
- [106] Ger Yang, David Basanta, and Georgios Piliouras. Bifurcation mechanism design—from optimal flat taxes to better cancer treatments. *Games*, 9(2):21, 2018.
- [107] Thiparat Chotibut, Fryderyk Falniowski, Michał Misiurewicz, and Georgios Piliouras. The route to chaos in routing games: When is price of anarchy too optimistic? *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [108] Georgios Piliouras. Catastrophe by design in population games: Destabilizing wasteful locked-in technologies. In *Web and Internet Economics: 16th International Conference (WINE)*, page 473, 2020.
- [109] Stefanos Leonardos and Georgios Piliouras. Exploration-exploitation in multi-agent learning: Catastrophe theory meets game theory. *arXiv preprint arXiv:2012.03083*, 2020.
- [110] Jakub Bielawski, Thiparat Chotibut, Fryderyk Falniowski, Grzegorz Kosiorowski, Michał Misiurewicz, and Georgios Piliouras. Follow-the-regularized-leader routes to chaos in routing games. *arXiv preprint arXiv:2102.07974*, 2021.
- [111] Yun Kuen Cheung and Georgios Piliouras. Vortices instead of equilibria in minmax optimization: Chaos and butterfly effects of online learning in zero-sum games. In *Conference on Learning Theory*, 2019.
- [112] Yun Kuen Cheung and Yixin Tao. Chaos of learning beyond zero-sum and coordination via game decompositions. In *International Conference on Learning Representations (ICLR)*, 2020.
- [113] Yuzuru Sato, Eizo Akiyama, and J Doyne Farmer. Chaos in learning a simple two-person game. *Proceedings of the National Academy of Sciences*, 99(7):4748–4751, 2002.
- [114] Ioannis Panageas and Georgios Piliouras. Average case performance of replicator dynamics in potential games via computing regions of attraction. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 703–720, 2016.
- [115] Boyu Zhang and Josef Hofbauer. Equilibrium selection via replicator dynamics in 2x2 coordination games. *International Journal of Game Theory*, 44(2):433–448, 2015.
- [116] Sai Ganesh Nagarajan, David Balduzzi, and Georgios Piliouras. From chaos to order: Symmetry and conservation laws in game dynamics. In *International Conference on Machine Learning (ICML)*, 2020.